

WHAT IS CLAIMED IS:

1           1. In a database system, a method for constructing an optimal query execution plan  
2 for executing a query, the method comprising:  
3           receiving a query specifying at least one join condition between two or more database  
4 tables;  
5           identifying each query block within said query, each query block comprising an  
6 atomic portion of said query;  
7           creating subplans for each query block based on grouping portions of each query  
8 block;  
9           determining at least one favorable access plan for each subplan of each query block,  
10 said at least one favorable access plan determined based at least in part on estimated  
11 execution costs;  
12           generating an optimal access plan for each query block based upon said at least one  
13 favorable access plan determined for each subplan; and  
14           constructing an optimal query execution plan based upon said optimal access plan  
15 generated for each query block.

1           2. The method of claim 1, wherein said step of identifying each query block within  
2 said query includes building a query optimization graph for each query block.

1           3. The method of claim 1, further comprising:  
2           generating a query optimization graph for each query block.

1           4. The method of claim 3, wherein said step of generating a query optimization  
2 graph includes generating subplans for each query block.

1           5. The method of claim 3, wherein said step of generating a query optimization  
2 graph includes generating plan nodes for each subplan, said plan nodes for joining tables and  
3 subplans.

1           6. The method of claim 5, wherein said step of generating a query optimization  
2 graph includes generating an array of access methods for each plan node.

1           7. The method of claim 5, wherein said step of generating a query optimization  
2 graph includes generating an array of join methods for each plan node.

1           8. The method of claim 1, wherein a subplan represents a table expression of a query  
2 block.

1           9. The method of claim 1, wherein said step of determining at least one favorable  
2 access plan for each subplan includes generating at least one access plan for each different  
3 set of outer references.

1           10. The method of claim 1, wherein the step of determining at least one favorable  
2 access plan for each subplan includes the substeps of:

3           placing a candidate plan segment in the next position in a current access plan being  
4 generated, said candidate plan segment representing a particular plan node, access method  
5 and join method valid at said next position;

6           evaluating the current access plan including said candidate plan segment; and

7           if the current access plan is less favorable than a favorable access plan previously  
8 identified, replacing said candidate plan segment with another available candidate plan  
9 segment and repeating said evaluating substep.

10

1           11. The method of claim 10, wherein a plan node comprises a quantifier object.

1           12. The method of claim 11, wherein said quantifier object represents a base table.

1           13. The method of claim 11, wherein said quantifier object represents a derived table.

1           14. The method of claim 10, wherein a plan node comprises a subplan object.

1           15. The method of claim 14, wherein said subplan object represents a join tree of a  
2 table expression of a query block.

1           16. The method of claim 10, further comprising:  
2           if the current access plan is more favorable than a favorable access plan previously  
3 identified, determining whether the current access plan comprises a complete plan;  
4           if the current access plan is determined not to comprise a complete plan, retaining the  
5 current access plan and repeating the above substeps for placing a candidate plan segment in  
6 the next position of the current access plan; and  
7           otherwise, if the current access plan is determined to be a complete plan, retaining the  
8 current access plan as a favorable access plan and repeating the above substeps to consider  
9 other available alternatives to the current access plan while alternatives are available.

1           17. The method of claim 10, further comprising:  
2           generating a property vector for said current access plan including said candidate plan  
3 segment.

1           18. The method of claim 17, wherein said property vector includes estimated  
2 execution costs for said current access plan.

1           19. The method of claim 17, wherein said property vector includes pipeline  
2 characteristics for said current access plan.

1           20. The method of claim 17, wherein said property vector includes order properties  
2 for said current access plan.

1           21. The method of claim 1, wherein said step of generating at least one favorable  
2 access plan for each subplan of each query block includes using a left-deep join enumeration  
3 strategy.

1           22. The method of claim 1, wherein said step of generating at least one favorable  
2 access plan for each subplan of each query block includes starting with the innermost nested  
3 subplans.

1           23. The method of claim 1, wherein said step of generating at least one favorable  
2       access plan for each subplan includes evaluating execution costs of partial access plans and  
3       pruning partial access plans less favorable than previously generated complete access plans.

1           24. A computer-readable medium having computer-executable instructions for  
2       performing the method of claim 1.

1           25. A downloadable set of computer-executable instructions for performing the  
2       method of claim 1.

1           26. In a database system, a method for generating a bushy trees during optimization  
2       of a database query, the method comprising:

3           receiving a database query specifying at least one join condition between two or more  
4       database tables;

5           identifying each query block within said query, said query block comprising an  
6       atomic block of said query;

7           building a query optimization graph for each query block, said query optimization  
8       graph including plan nodes representing subplans and quantifiers of each query block;

9           constructing a join tree for each subplan based upon selecting access methods, join  
10       methods, and join order for plan nodes of said query optimization graph having favorable  
11       execution costs;

12          constructing an optimal access plan for each query block based upon said join tree  
13       constructed for each subplan; and

14          generating a bushy execution tree based upon the optimal access plan determined for  
15       each query block.

1           27. The method of claim 26, wherein a query block comprises a selected one of a  
2       main block of a Structured Query Language (SQL) statement, a main block of a derived  
3       table, a main block of a view, a main block of a subquery used in a SQL statement, a derived  
4       table, and a view.

1           28. The method of claim 26, wherein said step of building said query optimization  
2 graph includes modeling left outer joins, right outer joins, and full outer joins as subplans  
3 which correspond to null-supplying sides of an outer join.

1           29. The method of claim 26, wherein said step of generating a query optimization  
2 graph includes generating an array of access methods for each plan node.

1           30. The method of claim 26, wherein said step of generating a query optimization  
2 graph includes generating an array of join methods for each plan node.

1           31. The method of claim 26, wherein said step of constructing bushy join trees for  
2 each subplan includes using a left-deep enumeration strategy.

1           32. The method of claim 26, wherein said step of constructing a join tree includes  
2 evaluating execution costs of a candidate plan segment to be added to said join tree being  
3 constructed.

1           33. The method of claim 32, wherein said candidate plan segment comprises a  
2 selected plan node together with an access method and a join method.

1           34. The method of claim 26, wherein said bushy tree comprises a processing tree  
2 having composite relations for left and right children of join nodes.

1           35. The method of claim 26, wherein said step of constructing an optimal bushy  
2 access plan includes evaluating execution costs of partial access plans enabling earlier  
3 pruning of unfavorable access plans.

1           36. In a database system, a method for optimizing execution of a query, the method  
2 comprising:  
3           receiving a query specifying selection of data from a plurality of database tables;

4 enumerating candidate plan segments for inclusion in an access plan for selecting data  
5 specified by the query, said candidate plan segments representing alternative strategies for  
6 joining relations and selecting data;

7 for each query block comprising an atomic portion of said query, determining an  
8 optimal access plan by performing the substeps of:

9 placing a candidate plan segment in a partial access plan being generated for  
10 said query block;

11 evaluating said partial access plan including said candidate plan segment;

12 if said partial access plan is less favorable than a complete access plan  
13 previously identified for said query block, pruning said candidate plan segment;

14 otherwise, adding an additional candidate plan segment to said partial access  
15 plan and repeating the above steps until a complete access plan for said query block is  
16 generated;

17 retaining a complete access plan if it is more favorable than other complete  
18 access plans previously generated for the query block; and

19 otherwise pruning a complete access plan which is less favorable than other  
20 complete access plans; and

21 generating a query execution plan based upon the optimal access plan determined for  
22 each query block.

1 37. The method of claim 36, wherein each said candidate plan segment comprises a  
2 plan node, an access method and a join method.

1 38. The method of claim 37, wherein a plan node comprises a subplan object.

1 39. The method of claim 37, wherein a plan node comprises a quantifier object.

1 40. The method of claim 36, where said substep of evaluating said partial access plan  
2 includes comparing estimated execution costs of said partial access plan to estimated  
3 execution costs of a complete access plan previously generated for said query block.

1           41. The method of claim 36, where said substep of evaluating said partial access plan  
2 before a complete access plan is generated enables early pruning of less favorable access  
3 plans.

1           42. The method of claim 36, wherein said step of generating an optimal access plan  
2 for each query block includes generating an optimal access plan for each subplan of each  
3 query block.

1           43. The method of claim 42, wherein said step of generating an optimal access plan  
2 for each subplan of each query block includes starting with the innermost nested subplans.

1           44. The method of claim 36, wherein said step of generating an optimal access plan  
2 for each query block includes using a left-deep join enumeration strategy.

1           45. The method of claim 36, further comprising:  
2 generating a property vector for each said partial access plan.

1           46. The method of claim 45, wherein said property vector includes estimated  
2 execution costs for a candidate plan segment.

1           47. The method of claim 45, wherein said evaluating substep includes comparing said  
2 property vector of said partial access plan to a property vector of a complete access plan  
3 previously generated for said query block.

1           48. The method of claim 45, further comprising:  
2 retaining a property vector of a complete plan generated for a query block if said  
3 complete plan is more favorable than any other complete access plan previously generated  
4 for the query block.

1           49. The method of claim 36, wherein said substep of pruning said candidate plan  
2 segment includes replacing said candidate plan segment with another available candidate  
3 plan segment and repeating said evaluating substep.